MODELING TRAFFIC PREDICTION UNCERTAINTY FOR TRAFFIC MANAGEMENT DECISION SUPPORT

Craig Wanke*, Sandeep Mulgund**, Daniel Greenbaum†, Lixia Song ‡

The MITRE Corporation
Center for Advanced Aviation System Development
7515 Colshire Drive, McLean, VA USA 22102

Abstract

Air Traffic Flow Management (TFM) is the process of balancing demand for airspace and airport resources with the capacity of those resources, in order to achieve both safe and efficient traffic throughput. Demand is typically estimated by predicting flight trajectories, and comparing the predictions to capacity metrics for airports and airspace. The effectiveness of TFM decision-making depends on the accuracy of these predictions. This effectiveness can be improved not only by improving prediction accuracy, but by quantifying the uncertainty in those predictions. When the uncertainty is known, decision analysis and risk management techniques can be applied to improve decision-making performance. To support this goal, a novel method has been developed for measuring and simulating uncertainty in traffic demand predictions. This method employs empirical observations of traffic characteristics to develop statistical models of the error distributions in demand predictions, which in turn can be used for Monte-Carlo simulation of specific traffic scenarios. Preliminary statistical results are presented here, as well as a discussion of simulation applications for both analysis and real-time decision-support tasks.

Background

Traffic flow management (TFM) is the process by which the Federal Aviation Administration (FAA), with the participation of airspace users, seeks to balance the capacity of airspace and airport resources with the demand for these resources. Together with the FAA's air traffic control (ATC) function, which provides for the safe separation of aircraft from each other and from restricted areas, TFM is a central component of the nation's air traffic management (ATM) system.

TFM personnel are known as Traffic Management Coordinators (TMCs) or Traffic Management Specialists (TMSs), depending on the facility in which they work. The general term for these personnel is *traffic managers*. One of their primary responsibilities is to ensure that traffic at national airspace system (NAS) resources (e.g., airspace sectors, airports) does not exceed levels that can be safely managed by controllers. Traffic managers also endeavor to ensure fair and equitable treatment for all NAS users, i.e., operators of commercial, general aviation, military, and other aircraft.

Traffic managers have many options when trying to address excess demand on a resource. For excess airport demand, a ground delay program is often used, in which arrival "slots" are rationed among airspace users, and flights are assigned delayed departure times such that available arrival capacity will be efficiently used. En route sector congestion, resulting from unusually high demand or when available airspace is limited due to hazardous weather, can be controlled several ways. Flights can be rerouted around hazardous weather and/or congested areas. Access to airspace can be limited by imposing miles-in-trail (MIT) restrictions at the airspace boundary, by applying ground delay, or in extreme cases by halting departures to some destinations (ground stop).

Decision support tools for TFM, therefore, must provide predictions of resource demand. Ideally, predictions should be provided based both on the current traffic situation and on proposed traffic management strategies, so that candidate solutions can be developed and compared. For example, the Enhanced Traffic Management System (ETMS)[1] used in the U.S. National Airspace System (NAS) provides real-time resource demand estimates based on predicted aircraft trajectories. In the near future, ETMS will be

*Senior Principal Engineer, Senior Member AIAA
**Lead Engineer, Senior Member AIAA
†Senior Engineer
‡Senior Engineer, Member AIAA

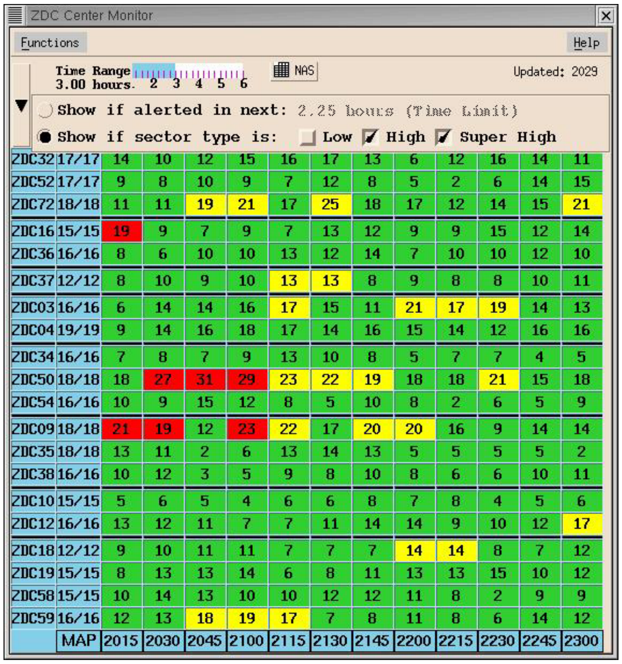

ZDC32	17/17	14	10	12	15	16	17	13	6	12	16	14	11
ZDC52	17/17	9	8	10	9	7	12	8	5	2	6	14	15
ZDC72	18/18	11	11	19	21	17	25	18	17	12	14	15	21
ZDC16	15/15	19	9	7	9	7	13	12	9	9	15	12	14
ZDC36	16/16	8	6	10	10	13	12	14	7	10	10	12	10
ZDC37	12/12	8	10	9	10	13	13	8	9	8	8	10	11
ZDC03	16/16	6	14	14	16	17	15	11	21	17	19	14	13
ZDC04	19/19	9	14	16	18	17	14	16	15	14	12	16	16
ZDC34	16/16	7	8	7	9	13	10	8	5	7	7	4	5
ZDC50	18/18	18	27	31	29	23	22	19	18	18	21	15	18
ZDC54	16/16	10	9	15	12	8	5	10	8	2	6	5	9
ZDC09	18/18	21	19	12	23	22	17	20	20	16	9	14	14
ZDC35	18/18	13	11	2	6	13	14	13	5	5	5	5	2
ZDC38	16/16	10	12	3	5	9	8	10	8	6	6	10	11
ZDC10	15/15	5	6	5	4	6	6	8	7	8	4	5	6
ZDC12	16/16	13	12	11	7	7	11	14	14	9	10	12	17
ZDC18	12/12	9	10	11	11	7	7	7	14	14	8	7	12
ZDC19	15/15	8	13	13	14	6	8	11	13	13	15	10	12
ZDC58	15/15	10	14	13	10	10	12	12	11	8	2	9	9
ZDC59	16/16	12	13	18	19	17	7	8	11	8	6	14	12
	MAP	2015	2030	2045	2100	2115	2130	2145	2200	2215	2230	2245	2300

Figure 1. ETMS Sector Count Monitor Display

capable of predicting resource demand as it would be affected by proposed reroute strategies[2], and research continues towards more sophisticated strategy impact assessment capabilities.[3,4]

Alerting of Excess Demand

The ETMS provides demand predictions for most NAS sectors in 15-minute bins, for prediction look-ahead times (LAT) of several hours. This information is available for particular sectors by user request, or in a collected form on a *Center Monitor* (CM) display as illustrated in Figure 1. The CM is a user-configurable display showing alerts for some or all of the sectors in a single Air Route Traffic Control Center (ARTCC).

Each cell in the CM matrix represents a 15-minute period, and the number in the cell represents the maximum predicted traffic count for any single minute within that 15-minute span. This value is often referred to as maximum instantaneous aircraft count (IAC) or simply "peak count" for the interval. The horizontal axis indicates increasing LAT (corresponding to 2015 to 2300 UTC, in this case). Each matrix row represents predictions for single sector (e.g. ZDC50). Next to the sector name are two sector alert thresholds (e.g. "18/18"), although currently, only one is used. This threshold is called the Monitor/Alert Parameter (MAP) and is compared to the peak count to determine whether a sector should be alerted. When the peak count is predicted to exceed the MAP for a sector, the corresponding box is colored yellow or red. Red alerts indicate that, of the aircraft involved in the peak count, enough are already airborne to exceed the MAP even if pre-departure flights are not counted. Otherwise, the alert will be yellow.

The MAP value is set to represent a traffic level high enough to be of concern to the traffic manager. It is not strictly accurate to refer to the MAP as a sector *capacity*, since there are many factors involved in sector workload beyond the number of aircraft present.[5,6] However, it is an easily-understood abstraction of workload for alerting purposes.

The alerts for all 20 centers can be aggregated on a single *NAS Monitor* display, which is similar in form to the CM except that each matrix row represents a whole ARTCC, and the numbers in the boxes reflect the number of alerted sectors in that ARTCC for the 15-minute time period. At the more detailed level, a single cell on the CM can be selected to produce a *Time-in-Sector* (TIS) display, which shows when individual flights will enter and exit a sector.

These sector traffic predictions are key TFM decision aids. Traffic managers use the alerts to identify areas of potential en route congestion, and by studying the flights and traffic flows involved, to identify candidate solutions such as reroute initiatives or MIT restrictions. Also, proposed TFM decision support systems[3] make direct use of these predictions when, for example, predicting the impact of a proposed reroute initiative. However, the usefulness of these predictions is a function of their accuracy. At the long LAT timeframes associated with strategic TFM decision-making, the predictions may not be very accurate.

Demand Prediction Uncertainty and Decision-Making

While traffic managers know that sector demand predictions are uncertain, they have very little information to use in quantifying that uncertainty and taking account for it when making decisions. ETMS sector load predictions include a crude estimate of uncertainty, in that alerts are differentiated into "red" and "yellow" based on whether or not all the aircraft involved are airborne. This is based on the assumption that departure time uncertainty is the largest source of uncertainty in the predictions. While this is useful for prioritizing traffic situations, in that the traffic manager would be justified in looking at red alerts before yellow alerts, it says little about the actual magnitude of uncertainty. For example, does a yellow alert of 3 aircraft over the MAP mean that there is an 80% chance that demand will exceed the MAP, or a 20% chance? Clearly, the answer to this question should influence the traffic management decision.

One way to factor in prediction uncertainty is to present probabilities directly or indirectly on traffic management decision support displays, relying on the skill of the traffic manager (and some procedural guidance) to use such information appropriately. Research is underway to understand the human factors issues in this area. Masalonis, *et al*[7] have developed candidate visualization methods for probabilistic sector demand information, in research that is directly linked to the work presented here.

Probabilistic predictions can also be used by decision support automation. Given detailed knowledge of demand prediction error distributions, standard decision analysis techniques can be applied to improve decision-making, assuming that a standard cost/utility criterion can be developed. For example, cost functions can be developed for both allowing excess demand and for taking action to prevent excess demand (e.g., rerouting flights). These functions would be partially subjective, reflecting the operational difficulties of managing excess demand or taking action to prevent it, and partially objective, reflecting the financial impact of traffic management actions on airspace users. The combination of these two functions would represent a *traffic management policy*. Applying decision analysis techniques and the known prediction uncertainties, decision support tools could compute "optimal" congestion management solutions with respect to this policy.

To use prediction uncertainty in any of these ways, it must be possible to measure and model it for a wide variety of specific traffic situations. The research presented here addresses this issue.

Previous Work on Demand Prediction Uncertainty

Previous work has been done on quantifying and modeling demand prediction uncertainty, using two primary approaches. Mueller[8] and Meyn[9] employed simulation approaches in which specific sources of trajectory prediction uncertainty were represented by closed-form statistical distributions, and Monte-Carlo simulation applied to determine the cumulative effects of the uncertainties on sector demand predictions. This technique is powerful, since it can theoretically be used to model arbitrary traffic and airspace situations. The component distributions can be controlled to simulate possible future changes in the operating environment (e.g., better data sources) and thereby evaluate the potential benefits of reducing the component uncertainties. It is feasible and instructive to use this approach for specific examples, such as in Ref. 8 where it is applied to a single sector, for a single aircraft type, on a single route. However, it is difficult and expensive to develop general purpose simulation models, with all required uncertainties, that work for a wide range of traffic types and conditions.

The second approach involves empirical study of the differences between predicted trajectories and actually-flown flight paths. This is subtle, because even for perfectly-predicted demand, the actual traffic may not correspond to the predictions. For example, if a large sector load excess is predicted, traffic managers will act to prevent it from occurring. However, the prediction of the *demand* may have been perfectly accurate, given the intentions of the flights at the time the prediction was made; the sector demand alert, in that case, worked precisely as intended.

The method used by Wanke, et. al.[10] overcomes this problem. By filtering out situations in which high demand levels were predicted, and traffic management actions were likely to have been taken, actual traffic counts could be compared to predicted counts to establish distributions of prediction error. The advantages of this method are that all uncertainties are captured in a small set of conditioned distributions (e.g., by primary sector traffic type), it is easy to apply to the entire NAS, and thus it provides valuable insight into the overall characteristics of prediction uncertainties in different kinds of airspace. The disadvantage is that by aggregating all uncertainty in a small set of distributions, it isn't possible to simulate different kinds or magnitudes of prediction uncertainties. It is also difficult to specialize the results to particular traffic situations, since the distributions reflect only peak counts, and cannot be applied to predictions of individual aircraft or specific flows.

Research Focus

The uncertainty modeling method described here attempts to address the deficiencies of both previously discussed methods. It is based on empirical study of the entire NAS, making it applicable to large-scale traffic analysis. It also breaks overall uncertainties down into categories whose parameters can be varied to simulate alternate prediction situations. Finally, it supports Monte-Carlo traffic simulations, so that uncertainty in real-world traffic situations can be modeled to support decision analysis studies and real-time decision support tools.

A Sector-Transit Approach to Modeling Prediction Uncertainty

Before discussing the modeling of prediction uncertainty, it is helpful to review the way trajectories are predicted by current ATC and TFM decision-support systems. Aircraft trajectories can be modeled

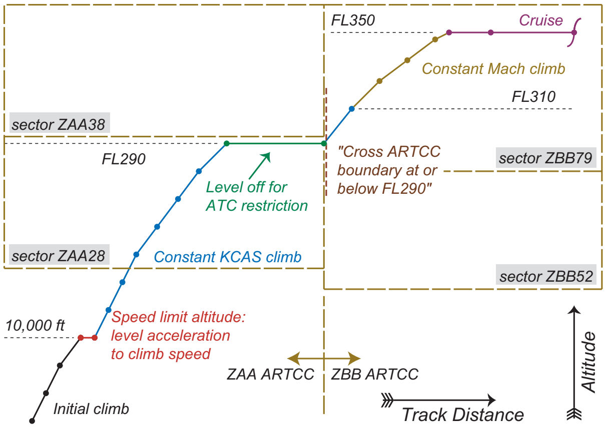

Figure 2. Trajectory Profile Prediction

at different levels of fidelity, depending on the application and the available input data. In air traffic management applications, the input data include filed flight plans, flight schedules and historical routes (for flights which have not yet filed plans), wind forecasts, track reports, and flight path constraints. The ETMS, and some other current systems and prototypes[3,11], use these data along with aircraft performance characteristics and a kinematic flight model to predict future positions. One decision-support system[12] uses a kinetic modeling approach in which the point-mass differential equations of aircraft motion are integrated. Regardless of approach, the resulting trajectory will contain a geographical (lateral) path based on the filed route, and a speed/altitude profile based on aircraft performance, winds, and ATC constraints. Figure 2 shows part of a fictional climb profile, to illustrate where prediction uncertainty arises.

The flight departs, climbing to the speed limit altitude of 10,000 feet MSL. In the U.S., flights must not exceed 250 knots of calibrated airspeed (KCAS) below this altitude. The aircraft accelerates to a climb speed, typically a constant value (such as 300 KCAS, for one typical four-engine turbofan aircraft) and resumes climbing. Frequently, the aircraft will need to level off briefly to satisfy an ATC constraint. These constraints make traffic flows more manageable for the controller teams. In this example, the constraint prevents climbing flights from entering sector ZAA38, which might be designed primarily to handle level traffic. These constraints are defined either in the Standard Operating Procedures (SOP) for an ATC facility, or in a Letter of Agreement (LOA) between adjacent facilities. Similar constraints may exist during the descent phase.

Several sources of uncertainty are apparent from this trajectory representation. For example, it is difficult to maintain accurate representations of aircraft performance for all equipment types in common use, and even were this possible, variations in parameters such as aircraft takeoff weight (unknown to TFM decision support automation) produce significant variations in actual performance. Wind forecasts contain uncertainty, and ATC constraints are not always applied precisely as stated in the LOAs and SOPs. Pilots may choose, for business reasons, to cruise at slower or faster speeds than filed in the flight plan. For trajectories of flights that have not yet departed, the departure time is difficult to predict precisely. The challenge addressed here is in quantifying these and other sources of uncertainty well enough to provide insight into, and improvements to, the TFM decision-making which is based on these predictions.

Trajectory Representation

The development of a model for estimating traffic demand uncertainty necessitated identifying a means for representing aircraft trajectories in the NAS. To support the objectives of this work, several requirements were identified for this representation. Specifically, it was determined that the approach should:

- Provide resolution of demand for NAS resources to the sector level
- Support human-in-the-loop (HITL) experiments in which subject matter experts are presented with probabilistic displays of sector demand
- Enable the generation of deterministic displays of the same traffic situations
- Allow generation of distributions and sensitivity analyses for TFM decision analysis studies
- Be computationally-tractable for generating large, statistically-valid trajectory populations

A crucial consideration was that decision support displays used in HITL experiments would present hypothetical situations that are consistent with the characteristics of real-world situations that traffic managers may encounter in the course of their work. To support the study of TFM decision analysis alternatives, it will be necessary to resolve traffic demand to the level of each individual aircraft within a time/airspace volume of interest. Earlier work[10] provided a simple measurement of overall uncertainty, but this would not be sufficient for the intended application here. These considerations motivated the examination of three possible trajectory representations:

- **Kinematic approximation**: This approach models the kinematic equations of motion of each aircraft in the time/space volume of interest using a series of constant-acceleration rhumb-line segments. The resultant representation is a piecewise linear approximation to a trajectory such as that shown earlier in Figure 2.

- **Directed sector transits**: With this approach, flight paths are abstracted as a series of time-stamped sector entries. By including the latitude/longitude coordinates of each sector boundary traversal, it is possible to reconstruct the general characteristics of an aircraft's flight from origin to destination. The segment between two successive sector boundary crossings is assumed to be traversed at constant-speed.

- **Cell transits**: Like the sector transit approach, this method abstracts an overall trajectory as a series of airspace cell traversals. Each cell provides a greater degree of spatial resolution than would a NAS sector definition.

Table 1 provides a comparison of each of these methods in relation to the requirements discussed above. For the intended application, it was determined that the cell transit method offered no significant benefits over the sector transit approach, but would explicitly require the definition of airspace cells, and their relationship to actual sector boundaries. While the kinematic approach would capture all of the details of an aircraft's flight between its origin and destination, it would be very computationally-intensive both to develop uncertainty models for every flight segment, and to employ them in the proposed Monte-Carlo simulation. It was therefore determined that the sector transit approach provided a good compromise.

Table 1. Candidate Trajectory Representation Methods

Method	Pros	Cons
Kinematic Trajectory	• Most widely applicable • Real routes/altitudes • Can adapt easily for future studies, traffic generation • Can vary distributions easily • Can calculate traffic complexity	• Difficult to model all required uncertainty distributions (e.g. for all aircraft types) • Computationally-intensive • Validation difficult • Must recompute sector crossings
Sector Transit	• Easiest to implement and validate • Can capture error sources readily, given available data sources • Facilitates modeling of unanticipated flights and reroutes • Due to aggregation, requires less data to develop distributions	• Harder to adapt to future studies, e.g. different NAS airspace organization • Statistics depend on sector map, which can change • Harder to model en route (cruise) speed errors than in kinematic representation • Difficult to vary conditions • Cannot easily evaluate traffic complexity within sectors
Cell Transit	• Similar advantages to sector approach • Discrete modeling techniques can be used • Applicable to solution search algorithms • Pre-computed sector map	• Difficult to validate • Difficult to gather statistical distributions • Requires formal definition of cell boundaries within existing NAS structure

Figure 3 illustrates the sector transit approach for representing aircraft trajectories. The dotted line represents a portion of an aircraft's trajectory through a hypothetical cluster of sectors, while the solid line represents the sector transit approximation that would be generated from this flight path. As shown, the aircraft climbs into high-altitude sector ZZZ06 from ZZZ02 and then levels off. Its subsequent lateral maneuvering illustrates the effects of the sector transit approximation. Specifically, the solid circles show the points that are used to build up the high-altitude sector transit list, which would consist of the following sectors in sequence: {ZZZ06, ZZZ27, ZZZ28, ZZZ20}. Each entry in the sector transit list contains the following information:

- Sector name
- Time of sector entry
- Altitude at sector entry
- Speed at sector entry
- Latitude/longitude at sector entry
- Flight phase (i.e., climb/cruise/descent/unknown)

This information can be used to reconstruct an approximation of an aircraft's flight from origin to destination. As shown in Figure 3, a result of this approach is that the details of an aircraft's maneuvering

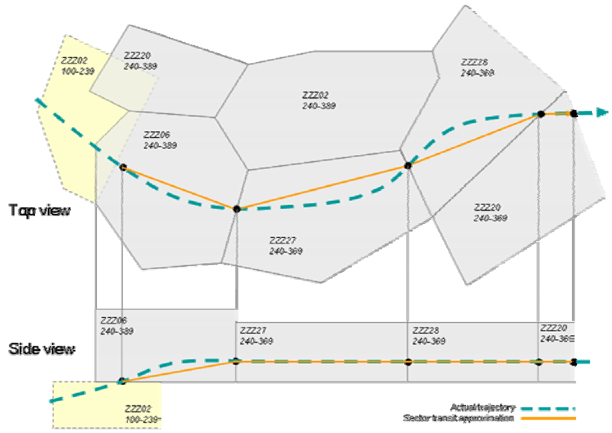

Figure 3. Directed Sector Transit Representation of Aircraft Trajectories

within a particular sector are lost, and its motion is treated as a straight line from one sector entry to the next. However, this is an adequate approximation for this effort's simulation requirements, which focuses on modeling prediction error distributions for TFM. Since the times and locations at which an aircraft enters and exits a sector are captured, this modeling approach captures the occupancy of individual NAS sectors as well as the overall flow patterns through them. The trajectory representation approach described above falls between physics-based simulations and direct measurement of the uncertainty.

<u>Uncertainty Modeling</u>

The purpose of building a library of sector transit representations of predicted demand and actual utilization of NAS resources is to develop insight into the distribution and sources of demand prediction errors. Specifically, given a nominal, deterministic prediction for a real-world traffic situation, what is the distribution of demand prediction errors for that situation?

In this approach, flights are treated individually, but trajectory prediction errors are grouped into observable error categories, which are assumed to be independent. This assumption is not guaranteed, and will be further discussed later. The error categories were chosen to maximize statistical independence and to span all major sources of error. There are six categories:

1) Flights that do not appear as expected (*no-shows*)
2) Unanticipated flights (*pop-ups*)
3) Routing
4) Altitude
5) Departure time
6) Flight progress

These are illustrated by Figures 4a-4d. In Figure 4a, a hypothetical airspace demand prediction is shown. Three flights are predicted, one of which is airborne at the time of the prediction; all flights pass through sector ZAA45. Figure 4b demonstrates error categories (1) and (2), which can be considered *flight existence* uncertainties. In this case, one of the predicted flights is cancelled (a no-show) and one new flight, which did not have an active flight plan at the time the prediction was made, appears (a pop-up). Note that these errors, as for all of the error categories discussed here, are conditioned on a particular prediction. For example, suppose prediction A is made at 1200 UTC for sector demand at 1400 (120 minute LAT), and prediction B is made at 1300. If an unscheduled flight, departing at 1330, files its first flight plan at 1230, then it will be considered a pop-up error for prediction A, since there is no information upon which to anticipate and model the flight at 1200. However, it will not be an error with respect to prediction B because at 1300 a flight plan is available, and thus a trajectory will have been predicted for that flight.

Figure 4c illustrates error category (3), in which a flight's routing is amended after the prediction was made, causing demand prediction errors for several sectors. Note that error category (4), in which a flight's altitude profile would be changed, would produce a similar effect. For example, the flight might travel through a low-altitude sector rather than a high-altitude one. Thus, categories (3) and (4) are *flight intent* uncertainties.

Figure 4d shows the impact of category (5) and (6) errors, both of which cause errors in predicting the position of the flight along its trajectory. These error types cause *flight progress* uncertainty. Previous studies have clearly demonstrated[13,14] that departure time prediction uncertainty is the single greatest prediction error component for pre-departure flights.

By categorizing the errors this way, distributions can be developed directly from empirical observation. Trajectory predictions, made at varying LAT, are compared to flight tracks, if the flight actually operated, to gain observations to develop distributions for each error category. These distributions form the basis of a Monte-Carlo simulation which can be used to estimate prediction uncertainty for a wide variety of traffic situations. The distributions themselves are also of interest, since they represent the relative contributions to overall uncertainty from the various error categories. Some preliminary results will be discussed later; first, the simulation process, source data selection and derivation of statistical error models will be described.

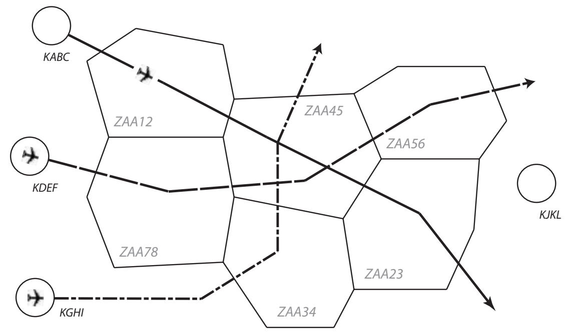

Figure 4a: Baseline Demand Predictions. Predictions for sector ZAA45 show three aircraft, one of which is airborne, to cross ZAA45 during the time period of interest.

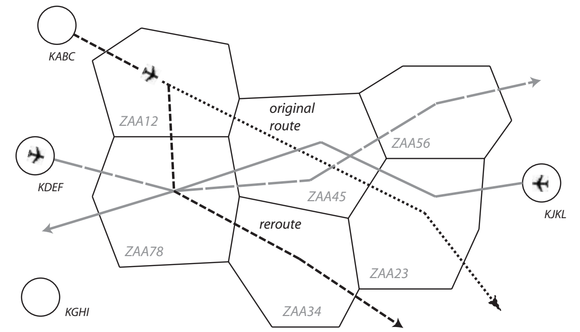

Figure 4c: Flight Intent Uncertainty. The flight from KABC is rerouted through ZAA78 and ZAA34 instead of ZAA45 and ZAA23.

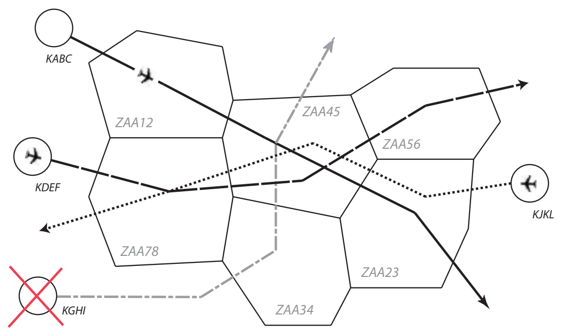

Figure 4b: Flight Existence Uncertainty. In reality, the flight from airport KGHI is cancelled, and an unexpected flight from KJKL appears.

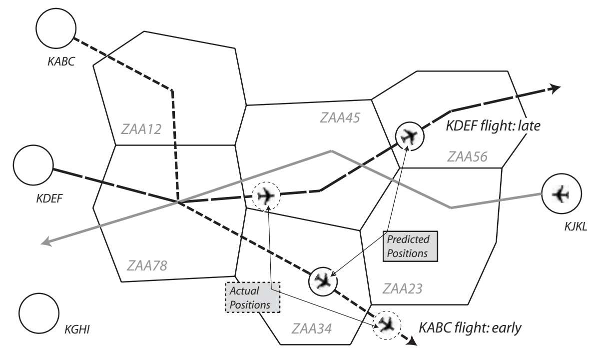

Figure 4d: Flight Progress Uncertainty. The flight from KDEF leaves 10 minutes later than predicted (departure time uncertainty) and the flight from KABC cruises at a higher speed than expected (sector transit time uncertainty).

<u>Monte-Carlo Simulation Process</u>

The fundamental simulation "question" can be posed as follows. Given a prediction time and a full set of NAS trajectory predictions for a range of look-ahead times (e.g. up to 6 hours ahead, for strategic TFM planning), what are the statistical properties of the prediction errors? Using the previously-defined error categories, the procedure for this calculation is as follows:

1. From the baseline trajectory predictions, compute the deterministic sector demand. This is analogous to how ETMS does this today.
2. Conduct a large number N of Monte-Carlo trials, each generating an alternate "actual" set of trajectories, for computation of stochastic sector demand:
   a. For each pre-departure trajectory in the baseline set, determine if the flight will be a no-show. If so, delete from the set for this run.
   b. For each trajectory, determine whether the flight intent will be changed (routing or altitude), and if so, select a modified trajectory for it.
   c. For each pre-departure trajectory, apply the departure prediction error distribution.
   d. For each trajectory, apply the sector transit time prediction error sequentially through the sector list.
   e. For each time bin in the prediction period, determine how many pop-up flights appear, select full trajectories for them, and add them to the trajectory set.
   f. Calculate sector demand from the modified trajectory set.
3. Once the N trials are complete, calculate desired demand statistics, by sector.

At this point, the sector demand uncertainties have been quantified. The results can be compared to the deterministic sector demand, used to drive probabilistic demand displays (as in Ref. 7), applied to decision analysis, or potentially used directly by real-time, automation-assisted decision support systems.

<u>Source Data Selection</u>

The raw data used to build the library of sector transits was obtained from ETMS. The full model will be based on transit data for the entire month of January 2003.

However, data samples presented later in this paper use a smaller set for preliminary analysis:

- Actual flight histories for the entire NAS between June 4 and 6, 2003, generated from ETMS's one-minute track reports
- Predictions in 15-minute intervals for the entire NAS between June 3 and 6, 2003. The predictions start earlier than the actual flight histories because any flights departing within the first six hours of June 4 would have associated ETMS predictions on June 3.

The analysis focused on en-route airspace under the control of an ARTCC. All flights not meeting these criteria were dropped from the data. The original raw text data was converted to an annotated extensible markup language (XML) format to facilitate analysis and integration with subsequent simulation applications. It was then necessary to "correlate" the data to match up each prediction with an associated actual flight, and to identify all of the predictions that were made for each actual flight record. This correlation process enabled the identification of pop-up flights and no-shows, since the former would be indicated by the presence of an actual trajectory with no prior predictions, and the latter would be indicated by one or more predictions but no actual flight record. Matching predictions with actual trajectories also facilitates identification of changes in flight intent, both before and during each aircraft's flight.

The process used for finding all of the predictions associated with an actual flight record was as follows:

- For each actual flight record, scan the relevant prediction data (up to six hours prior to the actual departure time).
- Identify those predictions whose ETMS flight index (an identifier that will be unique for that flight on that day) matches the index on the associated actual flight record.
- When a matching prediction is found, append the actuals flight record to indicate the timestamp on the correlated prediction. This provides a mechanism for subsequent retrieval of that prediction.

The converse process, for identifying the actual flight record associated with each prediction, was as follows:

- For each prediction record, scan the relevant actual flight data (up to six hours beyond the time at which the prediction was made)
- Determine if any of the actual flights have an ETMS flight index identical to that of the prediction. If the actual flight is on the next calendar day, also check for a match in flight ID (e.g., UAL1444).
- When a matching actual flight is found, append the prediction record to indicate the actual departure date of the flight associated with that prediction.

By performing this pairwise matching process across the entire set of ETMS predictions and actual flights, it was possible to construct a cross-correlated library for prediction error analysis.

Deriving Prediction Error Distributions for Simulation

Once correlated predictions and actual flight paths have been produced in sector-transit format, error distributions can be developed to support the Monte-Carlo simulation described earlier. The statistical model formulation for each required uncertainty category is described in the following sections. This formulation includes selecting the conditioning variables (CVs), collecting the observations, choosing a representation for the statistical distributions, and either storing the distribution directly (for sample-based distributions) or fitting parameters to closed-form statistical models.

It should be noted that the process described here, while generally applicable to any traffic or weather situation, is dependent on the chosen source data. In the initial version of the simulation, the distributions will represent a nominal statistical characterization of demand predictions for the NAS as of January 2003, given relatively good (i.e., no convective) weather. It would not be directly applicable to unusual traffic situations, such as when a large line of thunderstorms is present. Such analyses would require selecting different source data, and possibly different statistical model formulations.

No-shows

No-shows require the simplest statistical model. For each pre-departure in the baseline prediction set, the model must provide the probability of the flight becoming a no-show. The model is parameterized by the following CVs:

- Baseline departure (runway leaving, or "wheels-up") time, in hourly bins
- LAT to predicted wheels-up time, in 15 minute bins
- Origin
- Destination
- Operation Category (General Aviation, Air Carrier, or Military)
- Operator (Airline, or not specified)

- Estimated time enroute (ETE), in 15 minute bins
- Aircraft Class (Piston, Turboprop, Jet)

Observations are taken by running through the correlated prediction records, identifying whether the predicted flight actually operated, and recording the CVs and observations. The distribution model is a simple probability of no-show for each combination of CVs.

Routing/Altitude

In contrast to no-shows, flight intent errors are the hardest uncertainty to model. There are many different reasons that both flight routing and altitude changes can be made, and these reasons may not be easily observable in the data. For example, the presence of clear-air turbulence may cause large numbers of pilot-requested altitude changes.

Nonetheless, it is possible to construct a statistical model that represents the likelihood and characteristics of flight intent changes at a sector-transit level. The central hypothesis of the model is as follows. For each sector a flight enters, there is a measurable probability that the actually-flown sector list will begin to deviate from the predicted sector list. Should this occur, there are a number of possible alternate lists that can be flown, and these can be parameterized and sampled from an observed, empirical distribution.

The observations are made by comparing the sector transit lists, in time order, between each predicted trajectory and the corresponding actual trajectory. Each time the next sector in the predicted trajectory list corresponds to the next actually-transited sector, a flight counter is incremented according to the CVs listed below; this count is necessary to identify the total sample size, for later calculation of the probability that flight intent will change. If the next sector does not correspond, then it is assumed that the flight intent changed in the current sector, and an observation is stored. The observation consists of the *entire, as-flown trajectory* from the current sector forward, and becomes part of an empirical distribution for sampling in step 2b of the simulation process. In essence, there is a simple conditional probability, in each sector that a flight is predicted to enter, that the flight will undergo a sufficiently-large change in routing or altitude profile to change the future sector list. If, in a Monte-Carlo simulation run, this is calculated to occur, then a new trajectory from that point will be selected at random from trajectories that were observed to change in that sector. These observations are classified by the following variables:

- Origin
- Destination
- Current Sector
- Time-of-day
- Aircraft class
- Operation category

Departure Time

Departure time prediction errors are observed by taking pairs of correlated predicted and actual trajectories, and subtracting the predicted runway departure (wheels-off) time from the actual wheels-off time. Thus, positive values indicate that the flight departed later than predicted. These are collected and classified via the following CVs:

- Predicted wheels-up time (hourly bins)
- LAT to predicted wheels-up time (15 minute bins)
- Origin
- Destination
- Operation category
- Operator
- ETE (15 minute bins)
- Aircraft class

Sector Transit Time

Airborne flight progress prediction errors are captured by modeling errors in predicting the time it takes a flight to transit a sector. However, the time it takes a flight to transition through a sector is dependent on the precise path taken, which has been abstracted away by the sector-transit trajectory representation (as shown in Figure 3). To overcome this difficulty, a “directed transit” model is used. In this model, flights are classified by *sector triplets*, in addition to other variables of interest. The sector triplet includes the previous, current, and next sector in the transit list. Thus, flights will be grouped together by approximate flow pattern, even if there are differences between the exact routes of flight. For example, in Figure 3 above, the sector transit prediction error for sector ZZZ27 would be governed by the error statistics of flights traversing the triplet {ZZZ06, ZZZ27, ZZZ28}. This technique allows collection of prediction error statistics directly from trajectories in the sector-transit representation. Note that a sector triplet may begin with the flight origin rather than a sector, or end with a flight destination, if the error is being studied for the first or last en route sector in the flight’s trajectory.

There is a second difficulty. It is clearly unacceptable to treat sector time prediction errors in successive sectors traversed by the same flight as statistically-independent. If a flight is cruising faster than expected, for example, then the sector transit time should be over-predicted for many successive sectors along the flight path. This effect is captured by including the error experienced in the immediately preceding sector as a conditioning variable for the error distribution.

The process is as follows. For each pair of predicted and actual trajectories, the sector transit lists are compared in time order. For each sector with corresponding triplet between the predicted and actual trajectories, the sector transit time error is calculated by subtracting the predicted sector transit time from the actual sector transit time. The observation is stored according to the following CVs:

- Sector triplet
- Time-of-day (hourly bins)
- Transit time error from previous sector (signed, one-minute bins)
- Aircraft class
- Aircraft type (e.g. B73S)
- Operation category
- ETE

A parameterized distribution of the following form has been derived for these errors… [At the time of this writing, this data has not yet been analyzed; it will be complete for the final paper]

Pop-ups

Pop-ups, in the simulation (step 2e), are not specific to sectors or airspace; they are entire trajectories that are not anticipated at the time that the prediction was made, and thus will affect the demand in many sectors. Therefore, the statistical model for pop-ups is not explicitly conditioned on sectors or sector triplets. The model includes distributions of the number of pop-ups that will depart at each LAT (in 15 minute bins), and a classified list of trajectories from which to choose pop-ups when they are deemed to occur. In other words, step 2e in a simulation run is done as follows:

- For each LAT bin (0-15 minutes, 15-30 minutes, etc.) calculate, from the observed distribution, how many pop-up flights will appear.
- For each bin, select the requisite number of flights from the trajectory list.
- Calculate a new runway departure time for each flight from a uniform distribution over the interval covered by the LAT bin, and shift all sector transit entry times to correspond.
- Add these flights to the trajectory list for this simulation run.

Calculating distributions for each LAT bin is not difficult, since as described in the data collection process, each actual record contains a list of LAT bins for which it does and does not have a corresponding prediction. Since a full set of pop-up trajectories are being saved from which to sample, we do not need many conditioning variables; only those which affect the pop-up frequency are required. The following have been chosen as significant variables:

- LAT (15-minute bins)
- time-of-day at departure (hourly bins)
- day-of-week

These are used to govern both the pop-up frequency distribution, and the pop-up trajectory "pool" from which samples can be taken. The frequency distribution is NAS-wide, not sector-specific. So, although sector and sector traffic type are not variables, the differing effects of pop-up traffic on sectors is captured by this model. For example, sectors with large proportions of unscheduled general aviation traffic would be expected to have a higher pop-up rate. Since the captured pop-up trajectories will have a larger proportion of general aviation traffic than the overall trajectory population, sectors with such traffic will, correctly, exhibit a higher rate of pop-ups. Another sector-based difference, between arrival and departure sectors, will be illustrated in the next section.

Measured Uncertainty Characteristics: Pop-ups

At this writing, no completed simulation results are available. However, the statistical modeling data itself is of interest, since it provides insight into demand prediction uncertainty. One significant uncertainty source is illustrated here: the impact of pop-up flights on sector demand predictions, and how it varies with LAT and sector traffic type.

Pop-up distributions have been captured for the June 4-6, 2003 data previously described, and their impact on two sectors in the ZDC (Washington, D.C.) ARTCC has been analyzed. Figure 5 shows the impact of pop-ups on a sector which primarily handles departing traffic, for predictions of LAT 30 to 45 minutes, during four hours of the day (0900-1300 UTC). Under these conditions, there is an 85% probability that no unexpected flights will appear in the actual demand, with much smaller probabilities that one or two flights will appear. A fitted geometric distribution is also shown.

In contrast, Figure 6 shows predictions for the same sector at 120-135 minute LAT. In this case, the most likely result is still that no pop-ups will appear, but the probability is only 42%. There are significant probabilities that one or more flights will appear, including a 2% probability that five pop-ups will show up. Put another way, there is nearly a 40% probability

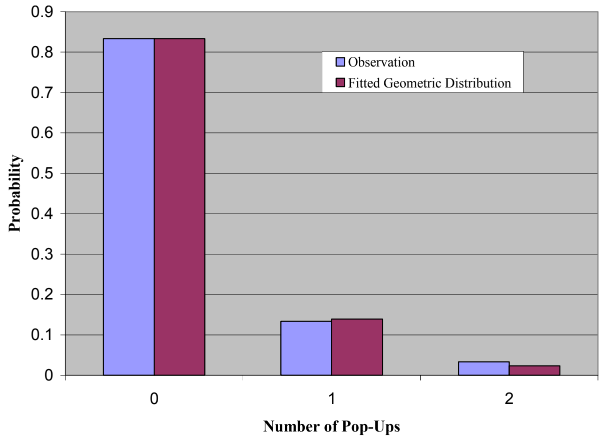

Figure 5: Pop-ups, Departure Sector ZDC05, 0900-1300 UTC, LAT 30-45 min.

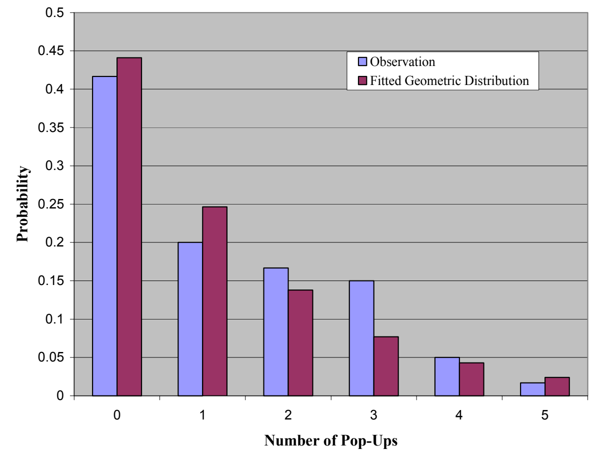

Figure 6: Pop-ups, Departure Sector ZDC05, 0900-1300 UTC, LAT 120-135 min.

that two flights will be present that were not predicted at all.

The situation looks different in a sector which primarily handles arrival traffic. Figure 7 illustrates the pop-up distribution for an arrival sector, also at 120-135 minute LAT. Compared with Figure 6, significantly fewer pop-up flights are expected to appear. This is consistent with the physics of the situation. In the departure sector, almost all of the flights predicted to enter the sector 120 minutes from prediction time will be on the ground, and in some cases will not even have filed a flight plan, at the prediction time. In contrast, flights arriving two hours from now will most likely have filed plans, and in many cases will already be airborne; hence, it is less likely that the decision support system will not yet know about them.

Previous work[10] has found that sector demand is, on average, under-predicted at long LAT. The pop-up distributions suggest one reason why this is so, though the other uncertainty sources also need to be studied to draw strong conclusions.

<u>Model Validation</u>

To date, the model has not been validated, since the statistical modeling process is not yet complete. Nonetheless, it is worth discussing what validation should be done, and what can be done with the available data sources. First, the model can be validated in an overall way, by comparing predicted

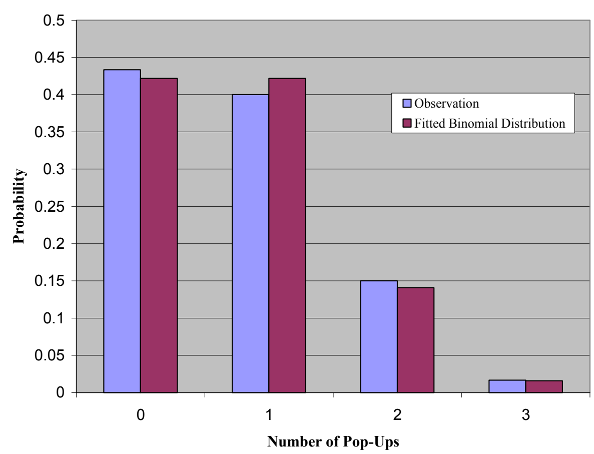

Figure 7: Pop-ups, Arrival Sector ZDC18, 0900-1300 UTC, LAT 120-135 min

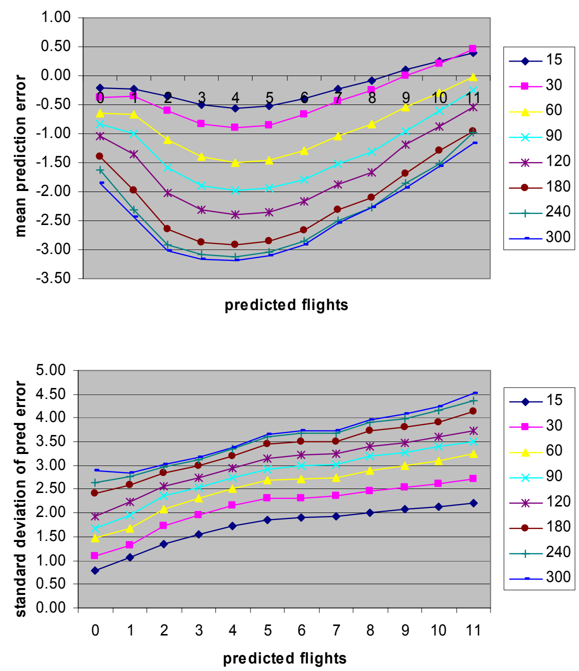

Figure 8: Mean and Std Dev. of Peak Count Prediction Errors for Mixed-Traffic Sectors by LAT (min)

sector demand uncertainties with measured ones. For example, see Figure 8, reproduced from Ref. 10. It shows the measured mean and standard deviation of errors in peak-count predictions in mixed-traffic-type sectors, as used by the ETMS Monitor/Alert function. These measured values are available for several sector types, at LAT ranging from 15 minutes to 6 hours. The results of the simulation described in this paper will be compared to these measurements as the first validation step.

There are several key assumptions in the model that require validation at a more detailed level. One such assumption is that the different uncertainty sources included in the model are statistically-independent. There are plausible reasons for this assumption to fail under some conditions. For example, if a scheduled air carrier flight departs late, it seems likely that it would cruise at a faster speed to make up time. Such correlations could be detected by statistical study of the joint distributions, and if warranted, additional variables can be added to the statistical models to account for such effects.

Another assumption was made in using sector triplets to classify sector transit time prediction errors. It is possible to evaluate this assumption directly by taking a sample of high-resolution trajectory data (i.e., not reduced to sector-transit form), evaluating the flight progress prediction errors, and then comparing the results to those achieved through the sector-triplet method. These, and other specific validation tasks, are planned for later in 2004.

Discussion

The simulation method described here has several important applications. The primary goal of developing it was to study TFM decision making. Once prediction error distributions are available, formal decision analysis techniques can be applied to study present-day decision making procedures, to develop better problem-solving strategies for common TFM problems, or even to develop new procedures that explicitly take prediction uncertainty into account when formulating solutions to a problem. Probabilistic demand displays are under development as a part of this work.[7]

It is also possible, given sufficient computing power, to use a simulation of this type for real-time decision support. In this application, the impact of a proposed TFM initiative could be evaluated not only in terms of the delay or workload that it would impose, but also in terms of the probability of success. This kind of feedback would allow traffic managers to decide, as one example, whether the initiative should be pursued immediately or delayed until the outcome is more certain. Finally, it is anticipated that TFM decision support tools will eventually be able to actively suggest solutions to traffic managers. It will be essential for such systems to measure and use demand prediction uncertainties – and possibly airspace capacity prediction uncertainties – when formulating suggested problem solutions. Knowledge of uncertainty also opens the door for solution strategies based on optimal estimation and control concepts.

Finally, aside from direct application to decision support research, there is merit in studying the magnitude and characteristics of prediction uncertainty in the NAS. The distributions captured for the simulation are themselves interesting, in that they provide a new look at the accuracy of present NAS demand predictions. It is possible, for example, to identify specific NAS sectors where predictions are particularly poor. This provides insight into how to improve predictions, either by improving modeling techniques, or perhaps by seeking new data sources which could reduce uncertainty.

Also, "increased predictability" is a possible benefit of many proposed ATC and TFM automation programs. Systems that support better airport surface traffic management might reduce departure time prediction uncertainty. Such effects can be modeled in the simulation by altering the departure time uncertainty model to reflect the proposed improvement. The simulation could then be run to evaluate the impact of the improvement on overall demand prediction uncertainty, and using a human-in-the-loop experiment or an automated set of decision rules, the decision-making benefits of the accuracy improvement could be evaluated.

References

1. *Enhanced Traffic Management System (ETMS) Reference Manual, Version 7.5*, Report No. VNTSC-DTS56-TMS-004, Volpe National Transportation Systems Center, U.S. Department of Transportation, November 2002.

2. Wanke, C, L. Carlson, J. DeArmon, S. Mayer, S. Miller, and G. Schultz, "New Tools for Strategic Route Planning: Collaborative Routing Coordination Tools," *Handbook of Airline Strategy,* edited by G. Butler and M. Keller, The McGraw-Hill Companies, New York, 2001.

3. Wanke, C., D. Berry, J. DeArmon, S. Miller, N. Taber, L. Rhodes, and M. Yee, "Decision Support for Complex Traffic Flow Management Actions: Integrated Impact Assessment," AIAA-2002-5015, *AIAA Guidance, Navigation, and Control*

Conference and Exhibit, Monterey, CA, 5-8 August 2002.

4. Sridhar, B., G. Chatterji, S. Grabbe, and K. Sheth, "Integration of Traffic Flow Management Decisions," AIAA-2002-3748, *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Monterey, CA, 5-8 August 2002.

5. Kopardekar, P., and S. Magyarits (2002). Dynamic Density: Measuring and Predicting Sector Complexity. *21st Digital Avionics Systems Conference*, Irvine, CA.

6. Masalonis, A., M. Callaham, and C. Wanke, "Dynamic Density and Complexity Metrics for Realtime Traffic Flow Management," *5th USA/Europe Air Traffic Management Research and Development Seminar,* Budapest, Hungary, 23-28 June 2003.

7. Masalonis, A., S. Mulgund, L. Song, C. Wanke, and S. Zobell, "Using Probabilistic Demand Predictions for Traffic Flow Management Decision Support," submitted to the *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Providence, RI, 16-19 August, 2004.

8. Mueller, K. Tysen, John A. Sorenson, and George J. Couluris, "Strategic Aircraft Trajectory Prediction Uncertainty and Statistical Sector Traffic Load Modeling," AIAA 2000-4765, *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Monterey, CA, USA, 5-8 August 2002.

9. Meyn, Larry A., "Probabilistic Methods for Air Traffic Demand Forecasting," AIAA 2000-4766, *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Monterey, CA, USA, 5-8 August 2002.

10. Wanke, C.,.M. Callaham, D. Greenbaum, A. Masalonis, "Measuring Uncertainty in Airspace Demand Predictions for Traffic Flow Management Applications," *2003 AIAA Guidance, Navigation, and Control Conference and Exhibit,* Austin, TX, 11-14 August 2003.

11. Brudnicki, Daniel J. and McFarland, Alvin L., "User Request Evaluation Tool (URET) Conflict Probe Results," *1st USA/Europe Air Traffic Management Research and Development Seminar,* Saclay, France, 17-20 June, 1997.

12. Chan, W., R. Bach, and J. Walton, "Improving and Validating CTAS Performance Models," *AIAA Guidance, Navigation, and Control Conference,* Denver, CO, August, 2000.

13. Krozel, Jimmy, Dan Rosman, and Shon R. Grabbe, "Analysis of En Route Sector Demand Error Sources," AIAA 2000-5016, *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Monterey, CA, USA, 5-8 August 2002.

14. Lindsay, K. L., D. Greenbaum, and C. Wanke, "Impact of Pre-Departure Uncertainty on Collaborative Routing Coordination Tools (CRCT) Prediction Performance," *AIAA 3rd Annual Aviation Technology, Integration, and Operations Technical Forum,* Denver, CO, USA, 17-19 Nov 2003.

<u>Notice</u>